Comment
Office of Science and Technology Policy
Request for Information on Artificial Intelligence

July 22, 2016

James C. Cooper*
Associate Professor of Law and Director, Program on Economics & Privacy
Antonin Scalia Law School
George Mason University

Introduction

Artificial intelligence and machine learning (collectively, "AI") offer great promise in myriad fields. As AI increasingly is used to make decisions of consequence, its potential to impacts consumers' lives will grow.

This comment is directed at the policy considerations related to AI decision-making in the commercial sphere. It argues that when determining whether any government intervention is needed to address such potential impacts on consumers from AI decision-making, policymakers should keep in mind two considerations. First, interventions should focus on instances that have been shown to be harmful, or are likely to be harmful. Second, policymakers should consider the extent to which market forces are likely to ameliorate any concerns; private incentives exist to correct biases, and they may operate more quickly and more efficiently than government action.

A. Harms-Based Approach

When developing policy, a starting point should always be a benefit-cost framework, centered on actual or likely harm. Policy should not be based on hypotheticals or remote risks.

An approach that focuses on actual or likely harms offers at least three advantages over an *ex ante* regulatory approach, which proscribes or mandates certain practices. First, by focusing on harm, one can be sure that government action is actually providing consumers with some benefits. Requiring harm to trigger action at least guarantees that the necessary (but not sufficient) condition for intervention to provide net consumer benefits is met.

Second, heterogeneous consumer preferences and costs of precaution increase the social costs associated with a common standard.[1] Third, a harm-based approach has the advantage of being more nimble than prescriptive rules, which would have to be reworked as technology or other market conditions change to alter the benefit-cost calculation with respect to certain AI practices. This consideration should weigh heavily, given the rapidly evolving nature of AI.

B. Identifying Harms

Policymakers should focus on conduct that causes, or is likely to cause, harm. One of the potential harms identified with respect to AI is biased decision-making. Before labeling the outcomes of AI decision-making harmful, however, one must be careful to take account of the context in which they occur.

One widely cited study, for example, found evidence that women were less likely than men to see an advertisement related to high paying jobs.[2] The computers trained to appear as women in the study instead were shown a generic job posting service. The inference taken by the authors is that "discrimination in the normative sense of the word" was at work, and that it could work to "further the current gender pay gap."[3] Such an inference, however, is based on incomplete information—one must consider the forces underlying the real-time auction to serve the particular ad.[4] A key factor driving the paper's findings is likely to be the willingness of other advertisers to bid to show ads to the female visitors. As competition for the attention of females increases, so will the price per impression, which could be too high for the employment ad. That is, the difference in ad serving rates is likely to be an artifact of more bidders competing for women's attention than men, rather than evidence of underlying bias in the AI system.

Similarly, another widely cited study reports differential online pricing for office supplies based on zip codes as evidence of a problematic algorithm.[5] In context, however, differential online pricing based on zip codes is common and much more likely related to heterogeneity in local costs and competition than underlying bias in an algorithm. A national chain has incentives to avoid having its ecommerce channel cannibalize its offline stores—which must respond to local supply and demand conditions. It accomplishes this procompetitive goal by equalizing offline and online prices to consumers across markets.

As these examples illustrate, policymakers should be hesitant to intervene when there is a beneficial (or benign) business explanation for the observed phenomena.

In addition to considering context, policymakers should focus on situations in which the output of classification from AI is likely to be harmful. Differences in online ad serving are unlikely to place severe constraints on opportunity sets, as they are only one source of information among many. For example, it is highly doubtful that a woman seeking employment will limit herself to opportunities presented in online advertisements on one particular web site.[6] Similarly, differential pricing is unlikely to pose problems for consumers, and in many cases is likely to increase consumer welfare, especially for relatively economically disadvantaged populations.[7]

C. Consider Market Forces

To the extent that firms are employing AI techniques that erroneously offer different commercial opportunities to different classes of consumers, policymakers must consider the competitive environment when deciding whether intervention is appropriate.[8] Private incentives exist to correct biases—biased AI decision-making that erroneously limits commercial options to certain disadvantaged populations represent a profit opportunity. For example, when a subprime automobile dealer was able to correctly distinguish systematic from transitory high-risk individuals, it was able to increase its profit while also increasing the amount of credit available to those who were relatively more credit worthy.[9] And competitive forces are likely to operate more quickly and more efficiently than government action to ameliorate such problems.

Relatedly, it is also crucial to distinguish between commercial and government use of AI decision-making. Unlike commercial entities that may use AI to incorrectly classify certain populations, governments are not subject to correction in the marketplace. Thus, governmental use of biased AI to make decisions regarding liberty or other fundamental rights, such as in criminal sentencing, are much harder to detect and correct.[10] Accordingly, governmental uses of AI should be the primary focus of policy concern.

Conclusion

AI offers great promise. As policymakers consider approaches to AI, they should focus on practices that are likely to be harmful to consumers, and ones that are unlikely to be corrected by market forces.

---

[*] This comment reflects the views of the author only. Affiliation is for identification purposes only.

[1] *See* James C. Cooper, *Separation, Pooling, and Big Data*, at 41-43 (April 2016), *at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2655794.

[2] Amit Datta *et al*., *Automated Experiments on Ad Privacy Settings*, 1 PROCEEDINGS ON PRIVACY ENHANCING TECHNOLOGIES, 92 (2015), *at* https://www.andrew.cmu.edu/user/danupam/dtd-pets15.pdf.

[3] *Id*. at 105.

[4] Indeed, the authors recognize as much. *See id*. at 93 ("Without knowledge of the internal workings of the [Internet advertising] ecosystem, we cannot assign responsibility for our findings to any single player within it nor rule out that they are unintended consequences of interactions between players.").

[5] Jennifer Valentino-Devries, Jeremy Singer-Vine, and Ashkan Soltani, *Websites Vary Deals and Prices Based on Users' Information*, WALL ST. J, (Dec. 12, 2012) (finding that differential online pricing based on zip code leads to those in relatively poorer zip codes to pay more).

[6] Indeed, the miniscule clickthrough rates—averaging around .1% in he U.S.—demonstrates how unlikely this scenario is. *See* David Chaffey, *Display Advertising Clickthrough Rates* (Apr. 26, 2016), *at* http://www.smartinsights.com/internet-advertising/internet-advertising-analytics/display-advertising-clickthrough-rates/.

[7] *See Executive office of the President, Differential Pricing* at 17 (Feb. 2015), *at* https://www.whitehouse.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf. The Antitrust Division has not brought a Robinson-Patman case since the 1960s, and the FTC has brought only one Robinson-Patman case since 1992. *See* ANTITRUST MODERNIZATION COMMISSION, REPORT & RECOMMENDATIONS at 318 (2007). Further, the bi-partisan Antitrust Modernization Commission recommended the repeal of the Robinson-Patman Act, concluding:

> [S]eventy years after passage of the Robinson-Patman Act, courts remain unable to reconcile the Act with the basic purpose of antitrust laws to protect competition and consumer welfare. . . There is no point in further efforts to reconcile the Act with the antitrust laws in general; the Robinson-Patman Act instead should be repealed.

Id. at 322.

[8] *See* FTC, BIG DATA: TOOL OF INCLUSION, Statement of Commissioner Maureen K. Ohlhausen, at A1-A2, (Jan. 2016), *at* https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf.

[9] *See* Liran Einav, Mark Jenkins & Jonathan Levin, *The Impact of Credit Scoring on Consumer Lending*, 44 RAND J. OF ECON. 249 (2013).

[10] *See* Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), *at* https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.